

Zeus User Meeting

April 4th, 2012

Agenda

- Overview of next weeks downtime(s)
- Upcoming changes to Moab
- Upcoming Changes to Modules
- Using MKL (Intel Math Kernel Library)
- Status of Restricted Data Protection on Zeus
- Running X applications from role accounts
- Known Issues
- Open Forum

Downtime Schedule

- April 10th – Facility Outage
 - Zeus and HPSS will be unavailable
 - Downtime will be start early (no later than 6AM EDT)
 - Downtime will last 8-10 hours
 - Facility work is expected to last about 2 hours (if all goes well)
 - If the facility work takes too long, we will restore the system and take a downtime on 4/11 to complete the system work.
- Additional system work during downtime
 - Install additional software on compute nodes
 - Reconfigure Moab
 - Remount Lustre filesystems to support protecting restricted data through access control lists (ACLs)
 - Additional system work
 - Change network switches, retest fabric

Upcoming changes to Moab

- During the downtime, changes will include
 - Block use of msub
 - Filter commands submitted via qsub to prevent certain options
 - Allow fallback to windfall to work properly
 - Get correct charging for batch, urgent, and windfall jobs

If you have been submitting jobs as we have documented (must specify a project with -A, do not use windfall) then you do not have to change anything.

Upcoming Changes to Modules

- From user experience and challenges, the module configuration will be changed
- This will affect very few users as it shouldn't change the order that modules load
- There will be no more compiler specific module names (netcdf-intel).
 - Compiler specific modules will be visible after the compiler is loaded
- The module system will provide consistent variables across all modules for specifying paths

If you are hard-coding paths into your makefiles, then it is a bug in our system. You should use the predefined variables.

Upcoming Changes to Modules

- Examples of how this will work

```
# module avail
-----/apps/modulefiles-----
Intel/11.1.072    Intel/12.0.4(default)        bbcp/10.0.4.46
Idl/8.1(default) nco/4.0.8(default)

# module load intel
# module avail
-----/apps/modulefiles-----
Intel/11.1.072    Intel/12.0.4(default)        bbcp/10.0.4.46
Idl/8.1(default) nco/4.0.8(default)
-----/apps/modulefamilies/intel-----
Netcdf/3.6.3(default) Netcdf/4.1.3
Hdf4/1.8.5(default)   hdf5/1.8.8(default)
Mpt/2.0.6(default)    impi/4.0.3(default)
```

Notice that new module options will be made available after the compiler is loaded.

Upcoming Changes to Modules

- There is some inconvenience to this, why change now?
 - We have learned a lot in the last 8 weeks
 - The sooner we do it, the easier the change
 - The changes bring good benefits
 - Minimize errors in linking conventions
 - Makes changing compiler families for development projects easier when “module netcdf” always does the right thing.
 - Cleans up the module listings

This change will happen in the next 30 days. We will give more information next week as to the timetable for the change, and how you can start testing the new modules. It will NOT happen during next week's downtime.

Using the Intel Math Kernel Library

- The Intel Math Kernel Library (MKL) provides many optimized math functions
 - BLAS
 - LAPACK
 - Scalapack
 - FFT
 - Many others
- Sometimes the linking convention can be tricky
- For basic math routings (BLAS, LAPACK, FFT), you can avoid the complexities by using
 - `-mkl=sequential`
- The MKL libraries are threaded, but since most all Zeus applications are already threaded with MPI or OpenMP, it is best to use the sequential versions of MKL

This information is also documented at:

https://nescdocs.rdhpcs.noaa.gov/wiki/index.php/Using_Intel_MKL

Using FFTs from MKL

- MKL provides a native optimized interface for FFT routines (DFTI)
- However, this interface is only portable to systems with MKL
- MKL also supports the FFTW3 (www.fftw.org)
 - “Fastest Fourier Transform in the West”
- FFTW3 is a optimized, open-source, library that runs on many different architectures
 - Intel, AMD, Power
- More importantly the FFTW3 interface is used by many vendors as an alternate interface.

Use of the FFTW3 interface brings portability across all NOAA machines, including the CCS.

Using Scalapack in MKL

- Scalapack is an MPI based library for scalable mathematical routines
- Linking is more complicated, because it is MPI library dependent
- MKL supports the major MPI stacks
 - MPT, Intel MPI, OpenMPI, MPICH
- To link on Zeus with MPT:

```
LDFLAGS=-L$(MKLROOT) -lmkl_scalapack_lp64 -lmkl_blacs_sgimpt_lp64 \
-lmkl_intel_lp64 -lmkl_sequential -lmkl_core
```

Protecting Restricted Data on Zeus

- Changes to get this to work will be made during the downtime on April 10th/11th
- No final plans as to how this will be done
- Meeting scheduled to discuss issue this week
- Why do we keep bringing it up?
 - Because status quo is going to change
 - Must protect the data
- Follow-up at the next scheduled meeting
 - Emails should be sent prior to that meeting

Running X applications from Role accounts

- While xsudo provides this support, it is a bit clunky and requires that an xterm be used.
- Xterms are slow due to X performance
- We have a new xsudo script that allows you to continue to use your shell window, but launch x applications like a normal user.
 - The feel is the exact same as a normal user
- The script will be updated during the downtime

Why can't we use ssh to login to a role account?

- Using ssh, even after logged into Zeus, is not as auditable as sudo.

Known Issues

- Known issues are documented at:
 - https://nesccdocs.rdhpcs.noaa.gov/wiki/index.php/Known_Issues



Any Questions?